

Cooperative Communication for High-Reliability Low-Latency Wireless Control

Vasuki Narasimha Swamy, Sahaana Suri, Paul Rigge, Matthew Weiner,
Gireeja Ranade, Anant Sahai, Borivoje Nikolić
University of California, Berkeley

Abstract—The Internet of Things envisions not only sensing but also actuation of numerous wirelessly connected devices. Seamless control with humans in the loop requires latencies on the order of a millisecond with very high reliabilities, paralleling the requirements for high-performance industrial control. Today’s practical wireless systems cannot meet these reliability and latency requirements, forcing the use of wired systems. This paper introduces a wireless communication protocol, dubbed “Occupy CoW,” based on cooperative communication among nodes in the network to build the diversity necessary for the target reliability. Simultaneous retransmission by many relays achieves this without significantly decreasing throughput or increasing latency.

The protocol is analyzed using the communication theoretic delay-limited-capacity framework and compared to baseline schemes that primarily exploit frequency diversity. In particular, we develop a novel “diversity meter” designed to measure “effective diversity” in the non-asymptotic regime. For a scenario inspired by an industrial printing application with 30 nodes in the control loop, total information throughput of 4.8 Mb/s, and cycle time under 2 ms, the protocol can robustly achieve a system probability of error better than 10^{-9} with nominal SNR below 5 dB.

Keywords—Cooperative communication, low-latency, high-reliability wireless, industrial control, diversity, Internet of Things

I. INTRODUCTION

The Internet of Things (IoT) vision is to enable a large number of globally distributed, embedded, computing devices to communicate with each other and interact with the physical world. This interaction includes not just sensing but also simultaneous actuation of numerous connected devices. For truly immersive applications, the latency requirements on the control loop are tens of milliseconds. This pushes the demand on the communication link latency to the order of a millisecond, while demanding very high-reliability. These requirements parallel those of modern industrial automation [1], with a round-trip delay of approximately 1 ms [2] and reliability of 10^{-8} [3], as achieved with wired connections.

This paper introduces “Occupy CoW¹,” a communication protocol for today’s industrial control and future IoT applications, designed to meet these stringent QoS requirements. The goal is to facilitate a plug-and-play transition from wired to wireless. This work builds crucially on [1] which established the need to attack this problem from the PHY/MAC layers and proposed a preliminary wireless architecture that focused on

low-latency operation through the use of reliable broadcasting, semi-fixed resource allocation, and low-rate coding. The main point of the present paper is that multi-user diversity can get us the desired reliability without relying on time or frequency diversity created by natural multipath or frequency selectivity. More importantly, as long as there are enough nodes present, this can be achieved robustly with small amounts of additional SNR. The result is largely independent of the model accuracy in the extreme tails of the fading distribution.

To motivate our protocol from the industrial control context, we first review the evolution of communication for industrial control and then briefly review cooperative communication and wireless diversity techniques. After that review, Section II describes our multi-user-diversity-based protocol in detail. Section III presents how it performs, how its internal parameters are optimized, and compares it to hypothetical frequency-diversity-based schemes. Section IV shows why the protocol is robust to uncertainty in fading models.

A. Industrial control

Communication in industrial control systems has traditionally been wired. Following trends in networking more broadly, point-to-point wired systems were replaced by *fieldbus* systems such as SERCOS, PROFIBUS and WorldFIP [5]–[7]. The main objective of fieldbus systems is to provide reliable real-time communication. There is a further desire to move to wireless communications for industrial control environments to reduce bulk and installation costs [8], and several wireless extensions of fieldbus systems have been examined [9], [10]. Unfortunately, these do not work in high-reliability settings since present designs for wireless fieldbuses are largely derivative of wireless designs for non-critical consumer applications and incorporate features such as CSMA or Aloha that can induce unbounded transmission delays [11]. On the other hand, ideas from wireless communication in Wireless Sensor Networks (WSNs) [12]–[14] that provide high-reliability monitoring also cannot be easily adapted for tight control loops because they tolerate large latencies [15].

The current generation of leading wireless technologies for industrial control are all based on successful WSN ideas. The Wireless Interface for Sensors and Actuators (WISA) [16] attempts to meet stringent real-time requirements, but fails to achieve interoperability and multi-path routing. The reliability of WISA (on the order of 10^{-4}) does not work for control [17]. ZigBee PRO [18] also fails to deliver high enough reliability [19]. Both ISA 100 [20] and WirelessHART [21] provide secure and reliable communication, but have relaxed latency bounds since they focus on non-time critical applications. These schemes are unable to hit the 2ms requirement we consider here. [19], [22]

¹OCCUPYCOW is an acronym for “Optimizing Cooperative Communication for Ultra-reliable Protocols Yoking Control Onto Wireless.” The name also evokes the similarity between our scheme and the “human microphone” implemented during the “Occupy Wall Street” movement [4].

There is a need for a faster and more reliable protocol if we want to have a drop-in replacement for existing wired fieldbuses like SERCOS III, which provide a reliability of 10^{-8} and latency of 1 ms.

B. Cooperative communication and multi-user diversity

Channel hopping, contention-based MACs and multi-path routing are commonly used time and frequency diversity techniques in WSN-inspired technologies [8]. However, most strategies for WSNs or industrial control networks do not exploit spatial diversity from multiple antennas or user cooperation, except implicitly through higher-layer approaches. Low-latency applications like ours cannot use time diversity since the cycle time is shorter than the coherence time. Techniques like Forward Error Correction and Automatic Repeat Request (ARQ) also do not provide much advantage [23]. Later in this paper, we demonstrate that frequency-diversity based techniques also fall short, especially when the required throughput pushes us to increase spectral efficiency. Consequently, our protocol leverages spatial diversity instead.

When there are multiple antennas in the system, we can harvest cooperative and multi-user diversity. Many researchers have studied these techniques in great detail; so our treatment here is limited. Laneman et al. [24] showed that cooperation amongst distributed antennas can provide full diversity without the need for physical arrays. [25] showed that even with a noisy inter-user channel, multi-user cooperation increases capacity and leads to achievable rates that are robust to channel variations. The prior work in cooperative communication tends to focus on the asymptotic regimes of high SNR and rely highly on the accuracy of the fading distribution in its analysis. By contrast, we are interested in moderate SNR regimes and do not want to rely on the accuracy of fading models. [26] found that reliable spectrum sensing by cognitive radios need not depend strongly on the details of fading distributions. It could instead rely upon the independence of fades across users. As we will see, it turns out the same idea applies here.

Multi-antenna techniques have been widely implemented in commercial wireless protocols like IEEE 802.11. [23], [27] use relays and a TDMA-based scheme to bring sender-diversity techniques to industrial control. Unfortunately, TDMA can scale badly with network size. To scale better with network size, our protocol uses simultaneous transmission by many relays, using some distributed space-time codes such as those in [28]–[30], so that each receiver can harvest a large diversity gain. This allows the protocol to achieve ultra-high-reliability without greatly decreasing throughput or increasing latency. While we do not discuss the specifics of space-time code implementation, recent work by Katabi *et al.* demonstrates that it is possible to implement schemes that harvest sender diversity using concurrent transmissions [31].

II. PROTOCOL DESIGN

The Occupy CoW protocol exploits multi-user diversity by using simultaneous relaying to enable ultra-reliable two-way communication between a central controller (C) and a set of n slave nodes (S) within a “cycle” of length T . Below is an abstract description of the protocol, which is also illustrated in Fig. 1. Distinct messages (size b bits) must flow in a star

topology from the central controller to individual nodes, and in the reverse direction from the slave nodes to the controller.

Following cellular convention, we refer to controller transmissions to slaves as downlink and slave transmissions to the controller as uplink. We assume that while normally, the controller and all nodes are in-range of each other, bad fading events can cause transmissions to fail. The protocol uses different nodes as relays to overcome this. On the downlink side, nodes that have received messages from the controller act as *simultaneous* relays to deliver messages to their destinations in a multi-hop fashion. A similar idea is applied for uplink. When they are not transmitting, all nodes are listening to everything. Nodes that have successfully decoded messages act as simultaneous relays for that message. This protocol is implemented by dividing every communication cycle into three phases each for downlink and uplink, with a small (but critical) scheduling and acknowledgment phase mixed in.

Resource assumptions

We make a few assumptions regarding the hardware and environment to focus on the conceptual framework of the protocol. All the nodes share a universal addressing scheme and order, and messages contain their destination address.

Fundamentally, errors are caused by bad fades. Since the short cycle time puts us in the non-ergodic flat-fading regime, time diversity cannot be used. All nodes are assumed to be capable of instantly decoding variable-rate transmissions [32]. All nodes are half-duplex but can switch instantly from transmit mode to receive mode.

Clocks on each of the nodes are perfectly synchronized in both time and frequency. This could be achieved by adapting techniques from [33]. Thus we can schedule time slots for specific nodes without any overhead. The protocol relies on time/frequency synchronization to achieve simultaneous retransmission of messages by multiple relays. We assume that if k relays simultaneously (with consciously introduced jitter²) transmit, then all receivers can extract signal diversity k .

A. Downlink and Uplink Phase I

Downlink Phase I (length T_{D_1}) is used by the controller to broadcast all messages to all slave nodes at rate $R_{D_1} = \frac{b \cdot n}{T_{D_1}}$ (Fig. 1, Column 1, nodes S0-S2 successfully receive the broadcast downlink packet). This is followed by Uplink Phase I (length T_{U_1}), in which the individual nodes transmit their messages (including one bit for an ACK/NAK to the downlink message) to the controller one by one according to a predetermined schedule at rate $R_{U_1} = \frac{b+1}{T_{U_1}/n} = \frac{(b+1) \cdot n}{T_{U_1}}$ by evenly dividing the time slots among all slave nodes. In Fig. 1, Column 2, nodes S0-S2 successfully transmit their uplink packets to the controller. Since all nodes are listening all the time, S2 and S0 are able to decode the message from S4, even though it does not reach the controller. Nodes successful in Phase I are called strong users.

B. Scheduling information

The scheduling phase (length T_S) is used by the controller to transmit acknowledgments to the strong users (Fig. 1,

²To transform spatial diversity into frequency-diversity [30].

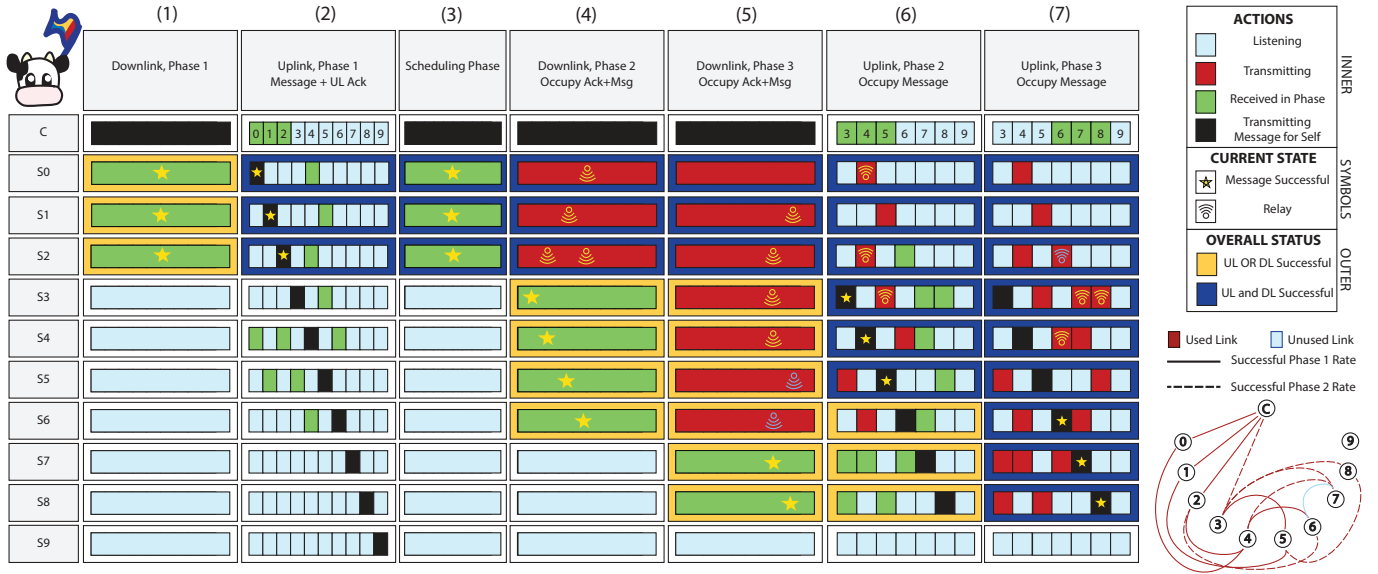


Fig. 1. The seven phases of the Occupy CoW protocol illustrated by a representative example. The table shows a variety of successful downlink and uplink transmissions using 0, 1 or 2 relays. S9 is unsuccessful for both downlink and uplink. The graph on the right shows the underlying link-strengths for the network.

Column 3). This is just 2 bits of information per slave node for downlink and uplink. This common-information about the system's state enables the controller and other nodes to share a common schedule for relaying messages for the remaining nodes. The strong nodes that are able to help must receive this information, and it doesn't matter that other nodes do not have this information at this time since they have nothing useful to say. This common-information is passed on to the remaining nodes in the downlink phase to follow.

The common ack information also allows the scheme to use possibly lower rates R_{D_2} and R_{U_2} , as we will see. The strong nodes S0-S2 in Fig. 1 receive the ack information.

C. Downlink Phase II and III

In Downlink Phase II (length T_{D_2}) the controller alters its broadcast message to remove already-successful messages for the strong nodes; so the packet is sent at an adapted rate, $R_{D_2} = \frac{n_1 \cdot (b+2)}{T_{D_2}}$, where n_1 is the number of nodes that were not successful in Phase I. In Fig. 1, Column 4, S3 receives the messages and ack information from the controller thanks to the lower adapted rate. The strong users (here S0-S2) compute the same adapted packet and act as simultaneous broadcast relays. Again, in Column 4, S4 receives the packet (including ack/scheduling information) through S2 and S0; even though the controller cannot directly reach it. S5 and S6 are similarly successful through relays S1 and S2, respectively.

Downlink Phase III (length T_{D_3}) follows the same structure as downlink Phase II, and transmits using the rate $R_{D_3} = \frac{n_1 \cdot (b+2)}{T_{D_3}}$. There exists the potential of three-hop relay paths from those who were successful in Phase II. For example, in Fig. 1, Column 5, S7 succeeds through S3. At the end of this phase, the nodes who received their messages from the controller have also received the global ack information. This allows these nodes to participate as relays in the uplink phases since they can calculate the uplink transmission schedule.

Note that the strong nodes that received the information from the controller in Phase I are the bottleneck for successful relay paths to other nodes during downlink.

D. Uplink Phase II and III

The calculated schedule from earlier phases allocates a slot for each unsuccessful node from Uplink Phase I in Phases II (length T_{U_2}) and III (length T_{U_3}). Time slots are again divided evenly among all n_1 unsuccessful slaves. In the slot for each failed slave node, the slave and everyone who heard that slave in an earlier uplink phase will simultaneously transmit the relevant message at the new rates $R_{U_2} = \frac{n_1 \cdot b}{T_{U_2}}$ and $R_{U_3} = \frac{n_1 \cdot b}{T_{U_3}}$.

This creates the potential for two-hop relaying if another slave heard the message in Uplink Phase I. For example, S2 and S0 transmit the message for S4 to the controller in Fig. 1, Column 6, since they already heard S4 in Phase I. Three-hop relaying is also possible in Uplink Phase III, for example the S6 \rightarrow S4 \rightarrow S2 \rightarrow C chain in Figure 1, Column 7. Note that this relies on S4 hearing S6 in Phase I, and S2 hearing S4 in Phase II. It is also possible to have new two-hop relay paths emerge due to the creation of new links (e.g. S7 to S3 in Phase II and S3 to controller in Phase III).

The uplink phases are similar to their downlink counterparts, but are in a sense inverted. The bottleneck to the controller now occurs on the last-hop, i.e. in Phase III.

As a final note, the exact transmission rates for each of the uplink and downlink phases depend on the time allocated and number of nodes remaining.

III. ANALYSIS OF OCCUPY CoW

We explore Occupy CoW with parameters in the neighborhood of a practical application, the industrial printer case described in [1]. Recall that the SERCOS III protocol [34] supports the printer's required cycle time of 2 ms with reliability of 10^{-8} . So we target a 10^{-9} probability of error for

Occupy CoW. The printer has 30 moving printing heads that move at speeds up to 3 m/s over distances of up to 10 m. Every 2 ms cycle, each head's actuator receives 20 bytes from the controller and each head's sensor transmits 20 bytes to the controller. If we assume access to a single 20MHz wireless channel, this 4.8 Mbit/sec throughput corresponds to an overall spectral efficiency of approximately 0.25 bits/sec/Hz.

A. Behavioral assumptions for analysis

We include the following behavioral assumptions in addition to the resource assumptions in Sec. II. We assume a fixed nominal SNR on all links with independent Rayleigh fading on each link. We assume a single tap channel³ (hence flat-fading) and because the cycle-time is so short, we use the delay-limited-capacity framework [35], [36]. We also assume channel reciprocity.

A link with fade h and bandwidth W is deemed good (thus no errors or erasures) if the rate of transmission R is less than or equal to the link's capacity $C = W \log(1 + |h|^2 SNR)$. Consequently, the probability of link failure is defined as $p_{link} = P(R > C) = 1 - \exp\left(-\frac{2^{R/W} - 1}{SNR}\right)$.

If there are k simultaneous transmissions⁴, then each receiving node harvests perfect sender diversity of k . For analysis purposes this is treated as k independent tries for communicating the message that only fails if all the tries fail.

We do not consider any dispersion-style finite-block-length effects on decoding (justified in spirit by [38]). A related assumption is that no transmission or decoding errors are undetected [39] — a corrupted packet can be identified⁵ and is then completely discarded.

B. Two hop downlink

In a 2-hop scheme, there are two shots at getting the message across. Failure is the event in which at least one of the n slave nodes has not received its message by the end. R_{D_1} is dictated T_{D_1} as described in II-A. Let the number of successful Phase I slaves be x_d . Then the rate of transmission in phase II, $R_{D_2} = R_{D_1} \cdot \frac{(n-x_d) \cdot T_{D_2}}{n \cdot T_{D_1}}$. Let the probability of link failure corresponding to R_{D_1} and R_{D_2} be p_1 and p_2 respectively. Then the probability that the controller to slave link fails in phase II given it failed in phase I is given by $p_c = \min\left(\frac{p_2}{p_1}, 1\right)$. The probability of 2-phase downlink system failure is thus:

$$P(fail) = \sum_{x_d=0}^{n-1} \binom{n}{x_d} (1-p_1)^{x_d} (p_1)^{n-x_d} (1 - (1-\tilde{q})^{n-x_d}) \quad (1)$$

where $\tilde{q} = p_2^{x_d} \cdot p_c$. Equations for other error probabilities corresponding to uplink and 3-hops are derived similarly and can be found in [40].

³Performance would improve if we reliably had more taps/diversity.

⁴We are ignoring a subtle effect here due to space limitations. The cyclic-delay-diversity space-time-coding schemes we envision effectively make the channel response longer. This pushes the PHY into the "wideband regime" in wireless communication theory, and a full analysis must account for the required increase in channel sounding by pilots to learn this channel [37]. We defer this issue to future work but preliminary results suggest that it will only add 2–3 dB to the SNRs required at reasonable network sizes.

⁵Consider all messages to include a 40 bit hash that is checked. This can be added to the underlying message size.




Fig. 2. The performance of Occupy CoW as compared with reference schemes for $b = 160$ bit messages and $n = 30$ nodes with 20MHz and a 2ms cycle time, aiming at 10^{-9} . The numbers next to the frequency-hopping scheme represent the amount of frequency diversity needed.

C. Results and comparison

Following [1] and the communication theoretic convention, we use the minimum SNR required to achieve 10^{-9} reliability as our metric to compare Occupy CoW to two other baseline schemes.

Fig. 2 looks at performance with fixed payload size $b = 160$ bits as the number of nodes, n varies. Initially the minimum required SNR for Occupy CoW decreases with increasing n , even through the throughput increases as $b \cdot n$, but the curves then flatten out⁶.

The topmost comparison scheme restricts uplink and downlink to the first hop of Occupy CoW. The required SNR shoots off the figure, because the throughput increases linearly with nodes, but there are no gains from multi-user diversity. The second scheme is purely hypothetical. It allows each node to use the entire 2 ms time slot for its own uplink and downlink message but without any relaying and thus also no diversity. This bounds what could possibly be achieved by using adaptive HARQ techniques.

The last reference curve represents a hypothetical (non-adaptive) frequency-hopping scheme that divides the bandwidth W into k sub-channels that are assumed to be independently faded. The curve is annotated with the optimal k . As k (and thus hops) increases, the available diversity increases, but the added message repetitions force the instantaneous link data rate higher. For low n the scheme prefers more frequency hops because of the diversity benefits. The SNR cost of doing this is not so high because the throughput is low enough that we are still in the linear-regime of channel capacity. For fewer than 7 nodes, this says that using frequency-hopping is great — as long as we can count on 20 or more independently faded sub-channels to repeat across.

It turns out that the aggregate throughput required (overall spectral efficiency considering all users) is the most important parameter for choosing the number of relay hops in our scheme. This is illustrated clearly in Fig. 3. This table shows

⁶This impact of multi-user diversity eventually gives way and the required SNR would start to increase for very large n .

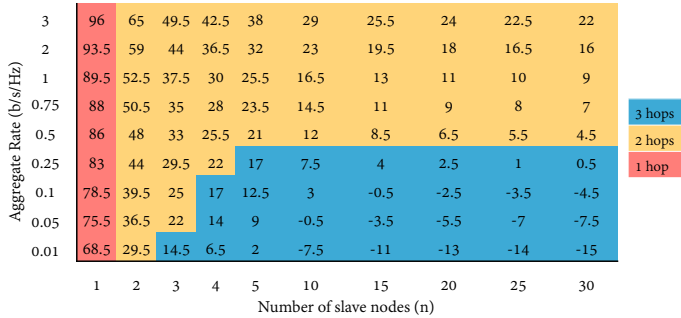


Fig. 3. The above figure tells us the number of hops and minimum SNR to be operating at to achieve a high-performance of 10^{-9} as aggregate rate and number of users are varied. Here, the time division within a cycle is unoptimized. Uplink and downlink have equal time, 2-hops has a 1:1 ratio across phases, and 3-hops has a 1:1:1 ratio for the 3 phases.




Fig. 4. Optimal fraction of time allocated for phase 1 in the 2 and 3-hop downlink as well as the 2 and 3-hop uplink schemes at various SNRs. Parameters used were 160 bit messages, 30 users, 2×10^4 total bits. Both downlink and uplink perform better with more time allocated for the bottleneck links to the controller; in downlink this means a longer phase 1, and in uplink this means a longer last phase.

the SNR required and the best number of hops to use for a given n . With one node, clearly a 1 phase scheme is all that is possible. As the number of nodes increases, we transition from 2-phase to 3-phase schemes being better. For $n \geq 5$, aggregate rate is what matters in choosing a scheme, since 3-phase schemes have to deal with a $3 \times$ increase in the instantaneous rate due to the shorter time, and this dominates the choice. In principle, at high enough aggregate rates, the one-hop scheme will be best even with more users. But when the target reliability is 10^{-9} , this is at absurdly high aggregate rates⁷. In the practical regime, diversity wins.

D. Protocol phase length optimization

It may seem natural to divide time evenly between phases so that links in different phases fail with the same probability, but Fig. 4 shows that performance can be improved with a slightly different allocation. It turns out that each phase of the protocol has a slightly different role to play in reaching the target reliability and this is seen by looking at the optimal phase lengths. For reasons of space, we do not discuss this here, and only point out the most important fact: for the kinds of SNRs that achieve system reliability 10^{-9} , the scheme prefers to spend slightly more time in the first phase. This is because the information about Phase I successes is shared

⁷We estimate this is around rate 40 — that would correspond to 40 users each of which wants to simultaneously achieve a spectral efficiency of 1.

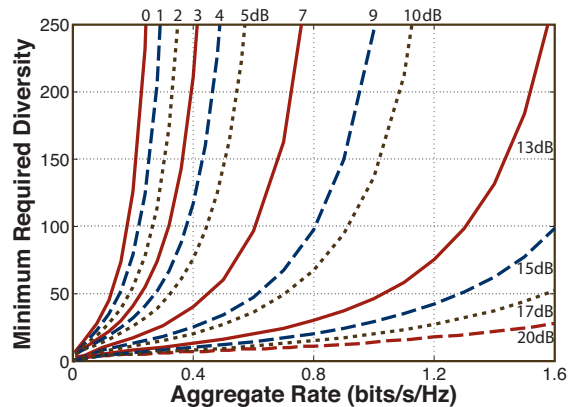


Fig. 5. Diversity Meter: Given an SNR, follow the appropriate line to the desired aggregate rate. The y-axis gives the minimal frequency diversity that a genie-aided frequency-hopping scheme would need in order to get $p_e = 10^{-9}$ at this SNR.

among strong nodes and allows for potentially lower rates in later phases. Since later phases tend to have fewer distinct messages to deal with, it makes sense to allocate slightly less time for them.

E. A “Diversity Meter”

Traditional formulations of diversity as a concept are asymptotic: they look at high SNR and very low error probabilities. While our scheme is also designed for the low probability of error regime, our SNRs are decidedly not asymptotically high. Studies of the finite-SNR diversity-multiplexing tradeoff tend to find that diversity is lower than expected in the non-asymptotic regime [41]. However, it is still unclear how the non-asymptotic diversity numbers should be interpreted. To resolve this question, we consider “frequency diversity” as the paradigmatic form of diversity, and propose a “diversity meter” as another baseline scheme. We use this to tell us how much frequency diversity the performance of a scheme like Occupy CoW is comparable to.

To get a fair bound, we do not want to restrict the diversity meter to non-adaptive repetition codes over frequency the way that we did earlier, but allow for intelligent coding. The frequency diversity is modeled as dividing the bandwidth into k independently faded blocks. An outer erasure code of rate $R_o = k'/k$ codes over all the blocks so the overall transmission succeeds if k' blocks are decoded. Each block uses a capacity-achieving inner code of rate R_i that succeeds if the block’s capacity is greater than R_i . The aggregate information rate is then given by $R_o \cdot R_i \cdot W$. The diversity required is defined as the smallest k , optimized over R_o and R_i to hit the desired performance. This is illustrated in Fig. 5 where one can look up the minimal frequency diversity required to achieve a probability of error of 10^{-9} at different combinations of aggregate rate and SNRs.

This turns out to be equivalent to schemes that know which of their diverse sub-channels are not too deeply faded and only transmit in those sub-channels (provided the encoding is not allowed to boost the local SNR in this channel to compensate for other deeply faded sub-channels).

The two-hop Occupy CoW scheme with 30 users and aggregate rate of 0.25 bits/s/Hz achieves a probability of

error of 10^{-9} at 0.42dB. With identical constraints on SNR, bandwidth and probability of error, Fig. 5 shows that the diversity meter requires a diversity of 201 with $R_o = 1/3$ and $R_i = 3/4$. Contrary to popular belief that we approach ideal diversity from below, this shows that at finite-SNR and for this intuitive sense of diversity, the multiuser diversity obtained with 30 users is comparable to having 201 independently faded sub-channels for every point-to-point link between controller and nodes! This drastic difference is because Occupy CoW gains diversity in one time slot instead of over many time slots, which allows for a lower effective R_i . Multiuser diversity is just better than frequency diversity in our context.

IV. ROBUSTNESS TO MODELING ERROR

Given the extremely low error probabilities we are targeting in a wireless setting, it is natural to question the impact of modeling error and uncertainty. Should anyone really trust the fading distribution down to 10^{-9} ?

To better understand this, let us consider the 2-hop downlink scheme with time evenly divided between both phases and no rate adaptation (i.e. all messages are repeated in both phases) $R_{D1} = R_{D2}$. For this case, the top curve in Fig. 6 shows the maximum probability of single link failure the system can tolerate for different number of nodes n , while keeping $P(fail)$ constant at 10^{-9} . The tolerable probability of link failure in the cases of interest ($n \geq 20$) is fairly high (above 15%), and fading models are quite good in this regime.

What if modeling error or the local industrial environment's interference introduced an extra probability of failure at each link, p_{env} , on top of the probability of error due to nominal fading, p_{fade} ? Then the probability of link failure is the combined effect of both these parts, p_{link} , is $p_{link} = p_{fade} + p_{env}$.

The bound on the tolerable p_{fade} can be obtained by shifting the p_{link} curve down by $p_{env} = 0.1$ in Fig. 6. The SNR labels show that to attain this new lower nominal link probability of failure for a given number of nodes, the SNR needs to increase. Fig. 7 this required SNR penalty increases for larger levels of robustness. But for really small network sizes ($n < 15$) achieving the probability of error requirement robustly can become impossible (when the desired tolerance to modeling error is itself greater than the maximum tolerable probability of link failure). For the moderate to large network sizes ($n \geq 25$), the protocol only has a small SNR penalty ($\approx 3dB$). We conclude that Occupy CoW does not rely on perfect knowledge of deep fading distributions and achieves high-reliability by relying on the independence⁸ of link failures.

Does the optimization of phase lengths from Sec. III-D greatly change in the presence of modeling errors? Fig. IV considers three different phase length allocations for the 2-hop downlink case: the naïve 50/50 allocation, an allocation optimized to account for both p_{fade} and p_{env} as part of p_{link} , and one that only chooses phase lengths based on p_{fade} (even though p_{link} contains p_{env}). We see that in the

⁸It is natural to wonder if we can rely on this independence. Preliminary results based on [42] indicate that even with a limited number of scatterers, multi-path will be essentially independent (or better) for the diversity perspective here. It is interesting to see that diversity contrasts with multiplexing gain — which is capped by the richness of the scattering environment.

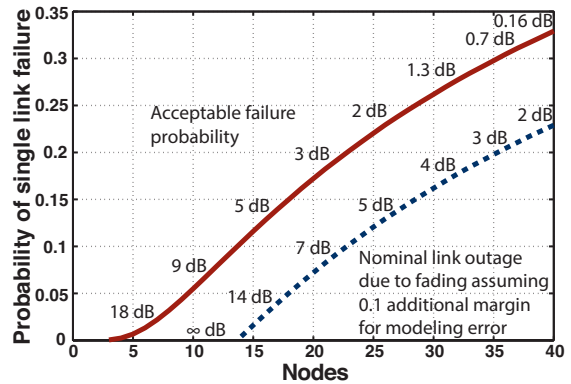


Fig. 6. The probability of link failure that can be tolerated in a 2-hop downlink as a function of the number of nodes. The aggregate rate is held constant at 0.25 bits/sec/Hz. The lower curve is 0.1 below and the SNR numbers represent the nominal SNR required to hit that particular probability of link failure for Rayleigh fading.

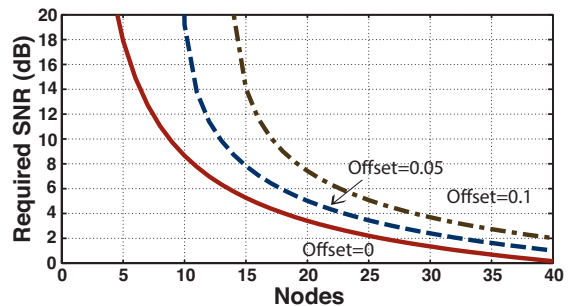


Fig. 7. SNR penalty to achieve performance robustly. For an increasing amount of robustness, a higher SNR is required as well as more nodes.

absence of modeling errors (i.e. accurate knowledge of deep fading distributions), the optimizing phase lengths provide a significant performance advantage over the naïve allocation. However, accounting for the modeling errors reduces the impact of optimization. This suggests that for phase length optimization to be effective in yielding a robust strategy, the fading model need not be highly accurate.

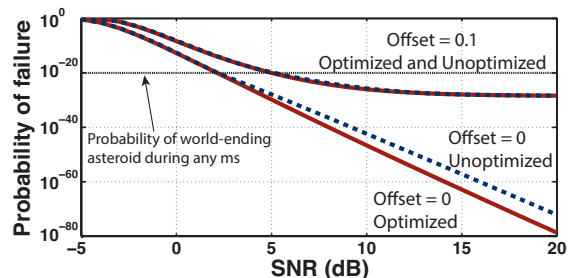


Fig. 8. Performance of three-hop downlink in the presence and absence of a worst-case modeling error that causes links to fail with a 0.1 higher probability regardless of SNR⁹.

⁹Approximate probability of catastrophic asteroid collision based on <http://neo.jpl.nasa.gov/risk/>. The last major asteroid event which wiped out the dinosaur population occurred 65 million years ago (approximately 10^{20} ms), leading to per millisecond probability of 10^{-20} .

ACKNOWLEDGEMENTS

Thanks to Venkat Anantharam, Milos Jorgovanovic and Kris Pister for useful discussions. We also thank the BWRC students, staff, faculty and industrial sponsors and the NSF for a Graduate Research Fellowship and grants CNS-0932410, CNS-1321155, and ECCS-1343398.

REFERENCES

- [1] M. Weiner *et al.*, "Design of a low-latency, high-reliability wireless communication system for control applications," in *IEEE International Conference on Communications, ICC 2014, Sydney, Australia, June 10-14, 2014*, 2014, pp. 3829–3835.
- [2] G. Fettweis, "The Tactile Internet: Applications and Challenges," *Vehicular Technology Magazine, IEEE*, vol. 9, no. 1, pp. 64–70, March 2014.
- [3] "SERCOS news, the automation bus magazine." [Online]. Available: http://www.sercos.com/literature/pdf/sercos_news_0114_en.pdf
- [4] E. V. Buskirk, "'Inhuman Microphone' app circumvents occupy wall street megaphone ban," 2011. [Online]. Available: <http://www.wired.com/2011/12/inhuman-microphone/>
- [5] R. Zurawski, *Industrial Communication Technology Handbook*. CRC Press, 2005.
- [6] S. K. Sen, *Fieldbus and Networking in Process Automation*. CRC Press, 2014.
- [7] A. Willig *et al.*, "Wireless Technology in Industrial Networks," in *Proceedings of the IEEE*, vol. 93, no. 6, June 2005, pp. 1130–1151.
- [8] P. Zand *et al.*, "Wireless Industrial Monitoring and Control Networks: The Journey So Far and the Road Ahead," *J. Sensor and Actuator Networks*, vol. 1, no. 2, pp. 123–152, 2012.
- [9] A. Willig, "An architecture for wireless extension of PROFIBUS," in *The 29th Annual Conference of the IEEE Industrial Electronics Society*, vol. 3, Nov 2003, pp. 2369–2375 Vol.3.
- [10] P. Morel *et al.*, "Requirements for wireless extensions of a FIP fieldbus," in *1996 IEEE Conference on Emerging Technologies and Factory Automation*, vol. 1, no. 1, Nov 1996, pp. 116–122 vol.1.
- [11] G. Cena *et al.*, "Hybrid wired/wireless networks for real-time communications," *Industrial Electronics Magazine, IEEE*, vol. 2, no. 1, pp. 8–20, Mar 2008.
- [12] I. F. Akyildiz *et al.*, "Wireless sensor networks: a survey," *Computer Networks*, vol. 38, no. 4, pp. 393–422, Mar 2002.
- [13] A. Bonivento *et al.*, "System Level Design for Clustered Wireless Sensor Networks," *IEEE Transactions on Industrial Informatics*, vol. 3, no. 3, pp. 202–214, Aug 2007.
- [14] M. A. Yigitel *et al.*, "QoS-aware MAC Protocols for Wireless Sensor Networks: A Survey," *Comput. Netw.*, vol. 55, no. 8, pp. 1982–2004, June 2011.
- [15] A. Willig, "Recent and Emerging Topics in Wireless Industrial Communications: A Selection," pp. 102–124, May 2007.
- [16] G. Scheible *et al.*, "Unplugged but connected [Design and implementation of a truly wireless real-time sensor/actuator interface]," *Industrial Electronics Magazine, IEEE*, vol. 1, no. 2, pp. 25–34, Summer 2007.
- [17] V. Gungor and G. Hancke, "Industrial Wireless Sensor Networks: Challenges, Design Principles, and Technical Approaches," *IEEE Transactions on Industrial Electronics*, vol. 56, no. 10, pp. 4258–4265, Oct 2009.
- [18] Z. A. Standard, "ZigBee PRO Specification," October 2007.
- [19] A. Kim *et al.*, "When HART goes wireless: Understanding and implementing the WirelessHART standard," in *IEEE International Conference on Emerging Technologies and Factory Automation*, 2008, pp. 899–907.
- [20] ISA100, "ISA100.11a, An update on the Process Automation Applications Wireless Standard," in *ISA Seminar, Orlando, Florida*, 2008.
- [21] International Electrotechnical Commission, *Industrial Communication Networks-Fieldbus Specifications, WirelessHART Communication Network and Communication Profile*. British Standards Institute, 2009.
- [22] J. Akerberg *et al.*, "Future research challenges in wireless sensor and actuator networks targeting industrial automation," in *Industrial Informatics (INDIN), 2011 9th IEEE International Conference on*. IEEE, 2011, pp. 410–415.
- [23] A. Willig, "How to exploit spatial diversity in wireless industrial networks," *Annual Reviews in Control*, vol. 32, no. 1, pp. 49 – 57, 2008.
- [24] J. Laneman *et al.*, "Cooperative diversity in wireless networks: Efficient protocols and outage behavior," *IEEE Transactions on Information Theory*, vol. 50, no. 12, pp. 3062–3080, Dec 2004.
- [25] A. Sendonaris *et al.*, "User cooperation diversity. Part I. System description," *IEEE Transactions on Communications*, vol. 51, no. 11, pp. 1927–1938, Nov 2003.
- [26] S. Mishra *et al.*, "Cooperative Sensing among Cognitive Radios," in *IEEE International Conference on Communications*, vol. 4, June 2006, pp. 1658–1663.
- [27] S. Girs *et al.*, "Increased reliability or reduced delay in wireless industrial networks using relaying and Luby codes," in *IEEE 18th Conference on Emerging Technologies Factory Automation, 2013*, Sept 2013, pp. 1–9.
- [28] F. Oggier *et al.*, "Perfect spacetime block codes," *IEEE Trans. Inform. Theory*, vol. 52, no. 9, pp. 3885–3902, 2006.
- [29] P. Elia *et al.*, "Perfect Space-Time Codes for Any Number of Antennas," *IEEE Trans. Inform. Theory*, vol. 53, no. 11, pp. 3853–3868, 2007.
- [30] G. Wu *et al.*, "Selective Random Cyclic Delay Diversity for HARQ in Cooperative Relay," in *IEEE Wireless Communications and Networking Conference (WCNC), 2010*, April 2010, pp. 1–6.
- [31] H. Rahul *et al.*, "SourceSync: A Distributed Wireless Architecture for Exploiting Sender Diversity," in *Proceedings of the ACM SIGCOMM 2010 Conference*, ser. SIGCOMM '10. New York, NY, USA: ACM, 2010, pp. 171–182.
- [32] S. Verdú and S. Shamai, "Variable-rate channel capacity," *IEEE Transactions on Information Theory*, vol. 56, no. 6, pp. 2651–2667, 2010.
- [33] Q. Huang *et al.*, "Practical Timing and Frequency Synchronization for OFDM-Based Cooperative Systems," *IEEE Transactions on Signal Processing*, vol. 58, no. 7, pp. 3706–3716, July 2010.
- [34] "Introduction to SERCOS III with industrial ethernet." [Online]. Available: <http://www.sercos.com/technology/sercos3.htm>
- [35] S. Hanly and D. Tse, "Multiaccess fading channels. II. Delay-limited capacities," *IEEE Transactions on Information Theory*, vol. 44, no. 7, pp. 2816–2831, Nov 1998.
- [36] L. Ozarow *et al.*, "Information theoretic considerations for cellular mobile radio," *IEEE Transactions on Vehicular Technology*, vol. 43, no. 2, pp. 359–378, May 1994.
- [37] A. Lozano and D. Porrat, "Non-peaky signals in wideband fading channels: Achievable bit rates and optimal bandwidth," *Wireless Communications, IEEE Transactions on*, vol. 11, no. 1, pp. 246–257, January 2012.
- [38] W. Yang *et al.*, "Quasi-Static Multiple-Antenna Fading Channels at Finite Blocklength," *IEEE Transactions on Information Theory*, vol. 60, no. 7, pp. 4232–4265, 2014.
- [39] G. D. Forney, "Exponential error bounds for erasure, list, and decision feedback schemes," *IEEE Transactions on Information Theory*, vol. 14, pp. 206–220, 1968.
- [40] Gireeja Ranade, Vasuki Narasimha Swamy, Sahaana Suri, Paul Rigge, "Analysis for Cooperative Communication for High-Reliability Low-Latency Wireless Control." [Online]. Available: <http://eecs.berkeley.edu/~gireeja/analysisOccupy.pdf>
- [41] R. Narasimhan *et al.*, "Finite-SNR diversity-multiplexing tradeoff of space-time codes," in *Proc. of IEEE Conference on Communications*, vol. 1, May 2005, pp. 458–462.
- [42] A. S. Y. Poon *et al.*, "Degrees of freedom in multiple-antenna channels: a signal space approach," *IEEE Transactions on Information Theory*, vol. 51, no. 2, pp. 523–536, Feb 2005.